

Precise™ for Web

User Guide

Version 10.0.x

IDERA

Precise™ for Web User Guide

Copyright © 2019 Precise Software Solutions, Inc. All rights reserved.

Document release version 1.0

Precise™, Precise Software™, the Precise™ Logo, Precise i³™, Precise Indepth™, Precise Insight™, Precise Savvy™, SmarTune™, Performance Warehouse™, Application Service Dashboard™, Precise for Storage Tiering™, Precise for Storage Tiering Plus Apps™, Precise for Database & Storage™, Precise for Applications™, Precise for Storage™, Precise Insight Inquire™, Performance Management Database™, and PMDB™ are trademarks or registered trademarks of Precise Software Solutions, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, decompilation and/or reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Precise Software Solutions, Inc. and its licensors, if any.

Certain third-party software may be distributed, embedded, or bundled with this product or recommended for use in connection with its installation and use. Such third-party software is separately licensed by its copyright holder. The list that includes the names of the copyright and license agreements can be found in the Release Notes document.

THE DOCUMENTATION IS PROVIDED "AS-IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PRECISE SOFTWARE SOLUTIONS, INC. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Global Headquarters

Brookhollow Central III
2950 North Loop Freeway West, Suite 700
Houston, Texas 77092
Phone: +1-877-693-1886
Fax: +1-650-898-1666

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Precise product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Customer Support

For technical assistance, visit our customer portal at <http://www.idera.com/support/productsupport> where you can find an extensive knowledge base, product updates and our online community forums. You can also contact our Customer Support Team using our customer portal, or go to <http://www.idera.com/support/productsupport> for a list of our support access numbers in your country.

Contents

- Introducing Precise for Web 8
 - About Precise for Web..... 8
 - Improved performance management..... 8
- Components of Precise for Web..... 8
 - Precise for Web browser agent..... 9
 - Precise for Web Dynamic Instrumentation agent 9
 - Precise for Web server agent..... 9
 - Insight Savvy for Network 9
 - Precise for Web Collector 9
 - Precise for Web FocalPoint 9
- How Precise for Web works..... 10
 - About improving your Web-based business 10
 - About streamlining your Web-based business..... 10
- What you can do with Precise for Web 11
- Where to get more information 11
- Precise for Web basics..... 13
 - About Precise for Web tabs..... 13
 - How most tabs are structured..... 13
 - About the Precise bar..... 14
 - About the Main area..... 14
 - About the Association area 14
 - About drilling down in context 14
 - About configuring Precise for Web settings 15

| | |
|--|----|
| About configuring Transactions & Users Settings | 15 |
| About configuring Display Settings | 18 |
| About configuring Tree View Settings | 18 |
| About configuring Time Frame Settings | 18 |
| Tasks common to most tabs | 18 |
| Switching to a different tab..... | 19 |
| Selecting a time frame | 19 |
| Selecting an entity..... | 19 |
| Sending an email message | 20 |
| Adding, viewing, and deleting Favorites | 20 |
| Determining which table columns to display | 20 |
| Copying data to the clipboard..... | 21 |
| Exporting to the Precise Custom Portal | 21 |
| Getting an overview of your application | 25 |
| About the Dashboard tab | 25 |
| How the Dashboard tab is structured | 25 |
| About the Main area in the Dashboard tab | 23 |
| About the Association area and Associated Findings area in the Dashboard tab..... | 23 |
| Examining performance over time | 24 |
| About the Activity workspace..... | 24 |
| How the Activity workspace is structured | 25 |
| About the Main area in the Activity workspace..... | 25 |
| About the Association area and Associated Findings area in the Activity workspace..... | 25 |
| How the Activity workspace is structured..... | 26 |
| Examining Web instance statistics..... | 32 |
| About the Statistics tab | 32 |
| How the Statistics tab is structured..... | 32 |
| About the Main area in the Statistics tab..... | 32 |
| About the Association area in the Statistics tab..... | 32 |

| | |
|---|----|
| About the entities you can examine..... | 29 |
| About the Tier entity..... | 29 |
| About the Instance entity | 30 |
| About the Performance Counter entity | 31 |
| How the Statistics tab can help you identify performance problems | 31 |
| Examining memory statistics | 31 |
| Examining CPU utilization statistics | 31 |
| Examining load..... | 32 |
| Web Findings | 33 |
| Frequent SLA Breaches | 33 |
| Working with the finding | 33 |
| Hotspot Detected..... | 33 |
| Working with the finding | 33 |
| Significant Backend Activity..... | 34 |
| Working with the finding | 34 |

Introducing Precise for Web

This section includes the following topics:

- [About Precise for Web](#)
- [Components of Precise for Web](#)
- [How Precise for Web works](#)
- [What you can do with Precise for Web](#)
- [Where to get more information](#)

About Precise for Web

Precise for Web provides you with the application performance management you need by proactively monitoring, diagnosing, detecting, and correcting performance problems before they become visible in your Web applications.

Addressing the needs of Web applications managers, system administrators, managers, quality assurance specialists and planners, Precise for Web provides a complete view of real user experience as well as Web application performance and resource consumption by capturing, measuring, and correlating performance metrics from all critical system components.

Precise for Web offers a wide range of data presentations to enable quick and accurate monitoring and performance analysis.

Recent operational data, for designated time periods, supplements operational data to enable operations staff to detect problematic conditions that occurred in the past and that are likely to occur again and again, if not corrected.

Performance data is achieved and a fully detailed display of the historical data is provided for use in performance, QA and capacity planning analysis.

Improved performance management

IT staff frequently search for answers to the following types of questions:

- Is a site compatible with the desired Service Level Agreement (SLA)?
- Are there performance problems for the Web application as perceived from the user browser and from the server side?
- If there is a performance problem, is it related to specific Web pages?
- If there is a performance problem, is it related to a specific URL and what Web page is impacted by the problem?
- If there is a performance problem, is it related to network performance?
- If there is a performance problem, is it related to the size of URLs that are part of the Web application and what Web page is impacted by the problem?
- If there is a performance problem, is it related to specific geographical locations?
- If there is a performance problem, is it related to specific users?
- Do users become impatient and abandon loading certain Web pages?
- Does the Web application generate request errors?
- Can usage patterns be associated with performance of the site from different locations?

Precise for Web can help you find the answers to these questions and more. In a regular Web application or e-business site poor user experience can lead to poor customer service, bad publicity, low work efficiency, overstaffing and increased employee costs. The Web application users experience is crucial to the success of your business.

Precise for Web can show you the user performance parameters captured at the end-user browser as well as the Web application performance parameters captured from the Web server side. It can sort them according to different categories and correlate them according to specific business transactions to provide you with an indication of their effect on business performance goals.

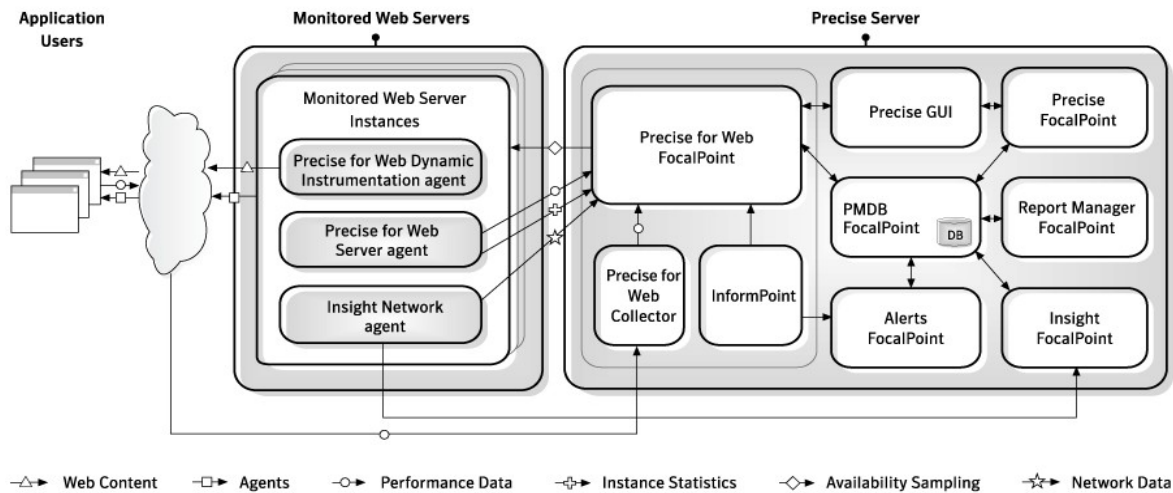
This helps you pinpoint where you should really invest your time and efforts into improving performance. Performance metrics by themselves are not enough. Understanding which performance problems are really affecting your business helps you to focus your resources into smart business decisions.

Components of Precise for Web

Precise for Web employs a client/server architecture that provides high performance Web application monitoring and tuning.

The figure below illustrates the typical architecture of Precise for Web.

Figure 1-1 Precise for Web architecture



Precise for Web consists of the following components:

- Precise for Web FocalPoint prepares the data for load and loads it into the PMDB. In addition the FocalPoint periodically samples Web server instances to check their availability.
- The end user's browser displays the monitored Web application. If Precise for Web is configured to collect client-side performance metrics, a browser agent which monitors the performance of the Web application from the client-side is loaded with the Web application pages.
- The Precise for Web Dynamic Instrumentation agent injects Java script callouts to the original Web content. The JavaScript is used to collect client-side performance data from the end-user browsers.
- A Precise for Web Server agent is installed on any monitored Web application instance that server-side performance metrics and statistics data is collected for.
- An Insight Savvy for Network is installed on the server of any monitored Web application instance that network performance metrics on the Web application instance is collected for.
- The Precise for Web Collector receives raw data from all of the Web server agents, the Web dynamic instrumentation agent, and the Network agent. The Precise for Web Collector correlates the different types of data and summarizes the data.
- PMDB.
- As part of the Precise product suite, Precise FocalPoint, Alerts FocalPoint, Report Manager FocalPoint, and Insight FocalPoint. Precise for Web is fully integrated with Precise products, allowing you to view Web application activity in Insight, receive alerts using Alerts, and examine Web applications activity reports within Report Manager. For more information, see the relevant product manual.

Precise for Web browser agent

The Precise for Web Browser agent, in the form of a JavaScript script, is installed on your Web pages and is sent to all end users. This agent collects Performance Data from the users' desktops and sends it to a Precise for Web Collector.

Precise for Web Dynamic Instrumentation agent

The Precise for Web Dynamic Instrumentation agent is responsible for injecting Javascript code, executed in end-user browsers, to collect client-side performance data. The Precise for Web Dynamic Instrumentation agent is composed of a filter engine which is installed on the monitored Web application instance and modifies the Web server response dynamically, just before it is sent to the browsers.

Precise for Web server agent

The Precise for Web Server agent is installed on the monitored Web application instance. The agent is a filter which is installed in the Web server and monitors URL activity and performance. In addition a process samples the Web server statistics and aggregates them into periods of 30-seconds. All the above data is transferred to the Web Collector using TCP protocol.

Insight Savvy for Network

Insight Savvy for Network is used to collect the network performance data of the Web application instance. For more information, see the relevant product manual. The data is transferred to the Web Collector for correlation with server-side data.

Precise for Web Collector

The Precise for Web Collector is installed on every machine that has one of the above-mentioned Web agents installed, including the Precise for Web FocalPoint.

The Precise for Web Collector:

- Receives all data collected by the agents on the machine it is installed on.
- Performs the basic processing.
- Correlates the different types of data and summarizes it.
- Creates files with the summarized data, that are later transferred to the FocalPoint machine for processing by the Data Loader.

Precise for Web FocalPoint

The Precise for Web FocalPoint is composed of two processes: a FocalPoint process and Data Loader process. The FocalPoint service handles several unrelated tasks:

- Answering data requests from the GUI or other Precise products.
- Cooperating with the PMDB to load the data files prepared by the Data Loader.
- The Data Loader loads collected data and prepares it for loading into the PMDB. The functionality of this process includes:
 - Sampling Web instances periodically to determine their availability.
 - Preparing data to be loaded into the PMDB.

How Precise for Web works

Browsers and Web application servers are the first and most common technologies used in today's Internet-enabled companies. The quality of a user's experience of a Web site, gauged by the response time, is a direct reflection of the company's commitment to customer relationship management and other similar business strategies.

Unfortunately, the business consequences of poor performance are greatest when the demand is at its peak. Precise for Web continually measures the business performance of a Web site by measuring key performance indicators and comparing those measurements against service level objectives.

Precise for Web provides detailed visibility into the performance of a Web application and the business transactions the application supports. The solution is leveraged to ensure the service provided to the Web application users is high-quality and the critical business opportunities supported using the Web application do not get disrupted.

Precise for Web provides customized data that can be collected from the client-side or from the server-side, such as the application user or any application title formatted in the Web page.

Precise for Web measures the real response time of end users to ensure the timely delivery of content to visitors. It tracks page views to determine the most active pages on the Web site, helping to prioritize areas for improvement.

Precise for Web measures the real response time of Web applications to provide pinpoint analysis of the problems that impact the end user. Using Precise for Web, you can prioritize problems not just by the fact that they occur but also by how severely they impact important business transactions.

Using Precise for Web, you can define and monitor service level thresholds to ensure customer satisfaction and retention. Precise for Web also tracks, abandons and requests errors to quickly determine potentially lost business and help prioritize redesign efforts.

Using Precise for Web, you can group users into geographic locations and connection types (high speed or dial-up) to provide the information needed to isolate the root cause of performance degradation. Also using Precise for Web, you can view the activity divided by the users that performed it.

About improving your Web-based business

Improving the quality of experience for Web application users requires the collection of key performance indicators (KPIs) to attract and maintain customers. Business models that leverage Web applications need to be continually monitored, and these KPIs are leveraged to ensure profitability and maximum return on technology investments.

Performance visibility over time is required to increase business efficiency and prioritize future investments. Knowing where to focus resources and the impact they will have on the bottom line provides the intelligence required to increase customer satisfaction and loyalty.

Providing a positive experience to users leads to higher work efficiency, shorter sales cycles, and ensures that business goals and user satisfaction goals are met.

Customer Relationship Management (CRM) using the Web depends heavily on a Web site's stickiness and its ability to provide self-service capabilities to customers. While this 24x7 customer interaction opens alternative business opportunities, it also creates challenges for ensuring business efficiency.

Metrics related to Web site performance can measure the success of funneling customers to this alternative support channel, but success can only be displayed if transaction volume and response time are among the metrics that are tracked.

Comprehensive visitor analysis is required to determine the effectiveness of Web-based content. Precise for Web allows you to track the behavior of individuals and groups, providing critical business information.

About streamlining your Web-based business

Focusing on the customer experience improves the relationship between the service provider and the user. Understanding the response time on a mission-critical business transaction provides a "report card" on how well customers are being serviced and their probability of retention.

Web sites will continue to play an increasing role in today's business strategy. Measuring KPIs of Web sites provides data to ensure faster, smarter decisions and maximum profits.

Recognizing inefficiencies is the first step in making improvements. Precise for Web delivers a methodology for identifying inefficient business processes as well as the ability to drill down and pinpoint the root cause of performance degradation.

What you can do with Precise for Web

Use Precise for Web to detect, diagnose, and resolve performance problems in your Web applications. Precise for Web helps you build performance metrics based on performance data gathered at an end-user desktop and/or at the Web server. The metrics can be correlated with the sites business goals to verify that improvements and related expenses are prioritized by their impact on the business.

Where to get more information

More information on Precise, its products, technical notes, and so on, can be found in the Release Notes document for this version.

Precise for Web basics

This section includes the following topics:

- [About Precise for Web tabs](#)
- [How most tabs are structured](#)
- [About drilling down in context](#)
- [About configuring Precise for Web settings](#)
- [Tasks common to most tabs](#)

About Precise for Web tabs

Precise for Web continuously collects high-quality metrics from the end-user's browsers and from the Web servers of your Web server-based applications and correlates these metrics to build a clear picture of application performance from the end-user perspective - response time.

The Precise for Web user interface is made up of three tabs. Each tab has a different focus, provides a different view into your Web application's performance, or identifies patterns in resource consumption.

The table below describes which tasks can be performed in each tab.

Table 2-1 Tabs in Precise for Web

| Tab | Description |
|------------|--|
| Dashboard | Lets you quickly identify which Web instance has a performance problem, providing the first glimpse into where the problem lies. |
| Activity | Lets you examine the application's performance over time and helps locate performance bottlenecks in its behavior. |
| Statistics | Lets you identify performance bottlenecks in the way Web instance resources are used. |

How most tabs are structured

Though each tab is structured differently, most tabs consist of two or three different areas. Each area can include different control elements, such as tabs and view controls, and displays information in various formats, such as tables, graphs, charts, or a tree. The third area, the Associated Findings area, contains findings. The various areas are related to each other in that performing an action in one area affects the information displayed in other areas on the page. For example, the Statistics tab contains an upper and a lower area. The lower area (Association area) shows information in a table format. Each row in the table represents an application. The upper area (Main area) displays general information on the selected entity as well as related applications.

The entities displayed in the Association area are associated with the selected entity or related application displayed in the Main area. At times, the relationship between the entity displayed in the Main area and those displayed in the Association area is that of parent to child, and sometimes it merely represents that there is a relationship between the selected entity and the entities displayed in the Association area (such as when a statement is the selected entity displayed in the Main area and the list of programs associated with the selected statement is displayed in the Association area).

It is possible to click a row and drill down to an additional level below the selected level or use the filter feature to determine which information is displayed in the table. When you perform an action on an element in the lower area (in this case the Association area), the information displayed in the upper area (in this case the Main area) changes to reflect the action you requested.

You can control the tab by doing any of the following:

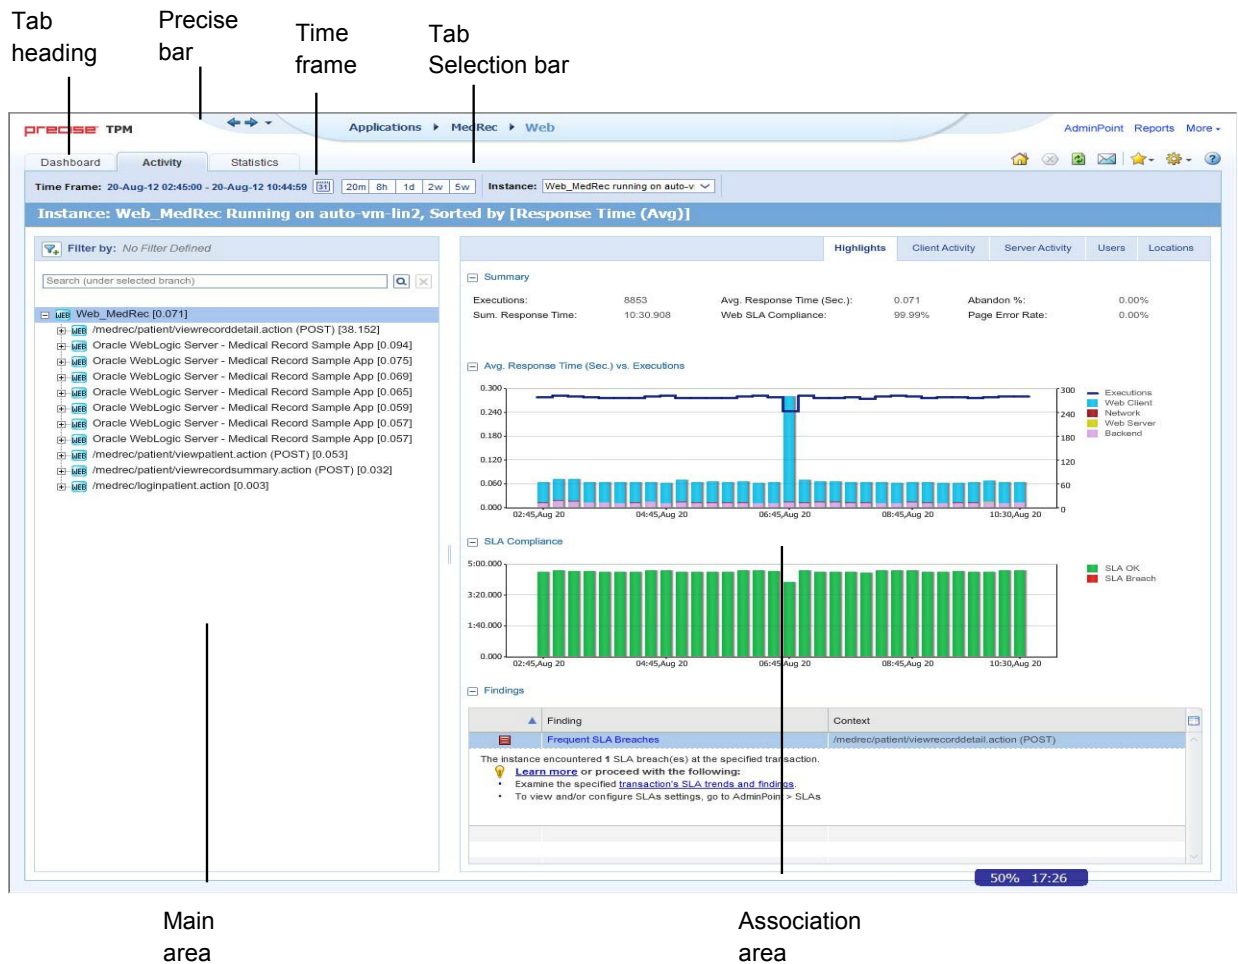
- Keeping the selected entity but changing the associated entities—by selecting a different entity from the Association controls in the Association area.
- Changing the selected entity—by selecting a row in the Association area.
- Changing the view for the selected entity—by selecting a different view in the Main area.

For example, when you want to view information on a specific application in the Current tab, choose Active Applications from the Association controls. The Association area changes to display applications-related information. The Tab heading and the Main area remain unchanged.

In the Association area, click the row of the application you want to view detailed information for. The Tab heading indicates the newly selected entity; the Main area displays information on the application you drilled down to, and the Association area shows information on statements executed by this application.

The following figure shows the structure of a typical tab.










Figure 2-1 How a typical tab is structured



About the Precise bar

The Precise bar enables you to keep track of where you have been and provides various controls. The following table describes the function of each of the toolbar buttons.

Table 2-2 Precise bar functions

| Icon | Name | Description |
|--|------------|---|
|  | Back | During a work session, keeps track of where you have navigated to. The Back and History buttons enable you to navigate between previously visited views. The Back control displays your previous view. |
|  | Forward | Enables you to navigate to the next view. This button is only enabled if you clicked Back or if you chose a history option. |
|  AdminPoint | AdminPoint | Launches Precise AdminPoint. |
|  | Home | Navigates to the highest level entity, usually the instance or Tier (all instances). The time frame settings remain the same. |
|  | Stop | Stops a request for information from the server. |
|  | Refresh | Updates the data currently displayed. |
|  | Favorites | Enables you to add or remove favorites in your Favorites list. |
|  | Send | Opens a new email message in your email program with the link to the current application in context. |
|  Settings ▾ | Settings | Opens the General Settings, Transactions & Users Settings, and Time Frame Settings dialog boxes. |
| Help | Help | Opens the online help in context. |

About the Main area

In most tabs, the left or top area is the Main area. This area displays an overview of all instances monitored by your Precise product. This information can be displayed in either graph, table, or tree format. The time frame for which the information is displayed is shown in the workspace toolbar. The times displayed are the local times on the FocalPoint server where the page is generated.

About the Association area

The Association area displays in-depth information for the entity selected in the Main area. Each workspace has analysis tabs specific for that workspace. As you navigate through Precise for Web, the analysis tabs change to enable you to view specific information relevant to the selected workspace and entity.

About drilling down in context

The term “in-context” means that you can display additional information on a selected item by drilling down to another workspace or view. The filter settings you define (for example, the time frame you chose) and the entity you select are carried over to the other view or workspace, to allow you to continue analyzing your subject from a different perspective. This concept takes on slightly different meanings depending upon where you are attempting to drill down in context from. Therefore, the information displayed when drilling down in context is always related to your original selection’s settings.

For example, when viewing information on an instance in the Dashboard workspace, you can click a link in the Details area (right pane) to view additional information on the related workspace, in the context of your original selection. Or when viewing a list of alerts for your product in Alerts, you can open your product in context and continue investigating the factors that led the system to issue that alert. This means that the information displayed when drilling down in context is always related to your original selection’s settings.

About configuring Precise for Web settings

The Transactions & Users Settings, Display Settings, Tree View Settings, and Time Frame Settings can be configured from the Settings menu on the Precise bar. They allow you to control the appearance and behavior of the user interface.

About configuring Transactions & Users Settings

On the Application User tab of the Transactions & Users Settings dialog box, you can configure the method to be used for collecting application users' data.

On the URL Parameters tab of the Transactions & Users Settings dialog box, you can configure the way the URL Parameters are collected and displayed.

Each settings tab displays an instances and clusters table that enables you to select a specific instance, cluster, multiple instances, or multiple clusters. By default, the instances and clusters table displays a list of all monitored Precise for Web instances and clusters that reside on the displayed Tier's node on the Tree View (where available). The user can control the way data is displayed in the table.

Click the triangle on the name column heading to sort instance and cluster data in ascending or descending order. Click on the filter icon displayed on selected table headings to filter data in that table column.

Table 2-3 Typical column headers on the tabs

| Heading | Description |
|--------------------|---|
| Name | Name of the instance and server or cluster. |
| Type | Type: Instance or cluster. |
| Application (Tier) | A list of the applications and Tiers that contain the instance or cluster. |
| Details | A short description of the instance or cluster current settings, according to the selected tab. |
| Server | Name of the server on which the instance or cluster is installed. |

Setting parameters at the cluster level affects all instances in that cluster. It is impossible to set parameters for an instance within a cluster individually.

The Cluster Instances button will be enabled when you select a cluster in the Transactions & Users Settings dialog box. Pressing this button will show the Cluster Instances dialog box. In this dialog box the instances of the cluster are displayed in a list.

Editing one or more instance settings

You can edit one or more instance settings. To edit one or more instance settings.

1. Click to select one or more instances for editing (for selecting more than one, press and hold down the CTRL key while clicking the instances you want to select).
2. Click **Edit**.
 Select the most suitable method to collect the application users' data. The two options are: Import settings or configure the user's collection method manually. The configuration format varies according to the selected method.
3. Click **OK**.

Editing one or more cluster settings

You can edit one or more cluster settings. For clusters, parameters are changed at the cluster level instead of the instance level. This means that a cluster parameter change affects all instances in that cluster. You cannot change the parameters of individual instances within a cluster.

To display a list of all instances in a cluster

1. In the Transactions & Users Settings dialog, select a Cluster row.
2. Click the **Cluster Instances** button. The Cluster Instances dialog is displayed listing all instances within the selected cluster.

To edit one or more cluster settings

1. Click to select one or more clusters for editing (for selecting more than one, press and hold down the CTRL key while clicking the clusters you want to select).
2. Click **Edit**.
 Select the most suitable method to collect the application users' data. The two options are: Import settings or configure the user's collection method manually. The configuration format varies according to the selected method.
3. Click **OK**.

Copy an existing user collection method configuration

The following procedure describes how to copy an existing user collection method configuration to the currently edited instance(s) and cluster(s). To copy an existing user collection method configuration

1. In the Edit Instance or Cluster dialog box, click **Import ...**
2. Use the application and/or Tier fields.
3. Click **OK** and the instance's or cluster's user collection method configuration will be copied to the Edit Instance or Cluster dialog box.

Manually configure the user collection method configuration

Manually configure the user's collection method by selecting a method in the 'Select method:' field.

NOTE The available methods in the 'Select method' field vary according to the selected instance(s) or cluster(s).

Table 2-4 Available methods for manual configuration

| Method | Pre-requisite instance configuration | Description |
|----------------------|---|---|
| Automatic | None | Try methods one after the other (until the application user is found) in the following order: <ol style="list-style-type: none"> a. Site Minder b. Packaged apps c. J2EE API or IIS API d. Remote User Variable e. Login Forms |
| User Defined Server | None | Allows extracting the user from the request header, cookie, or URL parameter. |
| Remote User Variable | None | Returns the REMOTE_USER Web server variable value. |
| SiteMinder | SiteMinder should be installed and configured to provide user information using the request header, cookie, or URL parameter (SM_USER). | Collect the user from the SM_USER key (request header, cookie, or URL parameter). |
| User Defined Client | Dynamic instrumentation agent installed. | Allows extracting the user from the cookie, URL parameter, or DOM Element. |
| Login Forms | Client-side collection (dynamic instrumentation) is installed. | Collect the user from HTML login forms. |
| PeopleSoft | PeopleSoft application and Client-side collection (dynamic instrumentation) is installed. | Collect the user from PeopleSoft HTML pages. |
| Oracle Applications | Oracle Applications and Client-side collection (dynamic instrumentation) is installed. | Collect the user from the Oracle Applications login form. |
| SAP Portal | SAP J2EE application and client side collection (dynamic instrumentation) is installed. | Collect the user from the SAP Portal login form. |
| IIS API | IIS Web server | Uses the information that IIS API provides. |
| J2EE API | J2EE Web servers (WebSphere, WebLogic, Tomcat, Oracle Applications server, SAP J2EE, SunOne) | Uses the information that J2EE API provides. |

About Login Forms configuration

When selecting the Login Forms option, the login forms table appears, which allows you to add/edit the list of URLs or forms, where the user name should be collected from. Clicking the Add or Edit button opens the dialog with the following fields (show table)

Table 2-5 Add / Edit fields

| Radio Button | Field | Description |
|--------------|------------|---|
| Login | Field Name | <p>The name of the username text box in the Web page. Consult the Web page administrator/developer, if applicable, to get this name. Alternatively you can find it by right clicking the mouse near the field and selecting "View source" from the popup menu. The structure of the HTML page should be like the following:</p> <pre><form name="loginForm"... action= ...> ... User Name: <input type="text" name="myUsernameField"> ... </form></pre> <p>The value you need to put in the "Field Name" field in this example is myUsernameField. Note that the input type is "text".</p> <p>Another way to get the field name: copy to the browser's address bar the following command, that should fill all the text boxes in the page with their name and the name of the form.</p> <pre>javascript:var frms = document.forms ;for (var formInd=0 ; formInd <frms.length ; formInd++) {var form = frms[formInd];var elems = form.elements;for (var ind=0; ind< elems.length; ind++){var field = elems(ind);if (field.type == "text")field.value="Field=" +field.name + " form=" + form.name;}};void(0);</pre> |
| | URI | <p>To save the time of looking for the field in all pages, specify the URI of the login page, e.g.</p> <pre>/login.jsp</pre> <p>TIP: When the login page opens in a dialog that has no address bar, switch to full screen mode to see its URI (F11 in most browsers).</p> |
| | Form Name | <p>If the URI is dynamically changing, or you can reach the login page using many URIs, you can mention instead (or in addition) the form name in which the user name field exists. You can find the form name in the same way mentioned above; in this example the value for the form name would be loginForm.</p> |
| Logout | URI | <p>Specify the URI of the logout page that after visiting makes sure that the user is no longer signed in , e.g. <code>/signOut.jsp</code>. All pages in a session since visiting the login page are marked as visited by the found user. Pages visited after this URI are no longer related to the user.</p> |

About manually configuring URL parameters

On the URL Parameters tab of the Transactions & Users Settings dialog box, you can define rules that map URLs that possess dynamic parameters to a format that can identify these Web pages and prevent them from having dynamic URLs.

In this procedure, you do the following:

- Select one or more instances or clusters in the table for which you want to configure the URL parameters.
- Define how to display the parameters in the parameter part of the URL.
- Define rules for specific parameters.

See [Editing URL mapping rules for an instance or cluster in the system](#) on page 18.

Editing URL mapping rules for an instance or cluster in the system

You can define how to handle parameters in the parameter part of a URL by using the URL Parameters dialog box. This dialog box also lets you specify exceptions to the defined rule.

NOTE A URL Mapping Rule is only applied to data that was collected after the rule was configured. Data that was collected before the rule was configured is not affected.

1. On the Precise bar, select Settings>Transactions & Users Settings.
2. Click the URL Parameters tab.
3. Select the instance or cluster whose URL mapping rules you want to modify and click **Edit**.
4. Define the system's default behavior when it encounters parameters in the parameters part of the URL.

Select one of the following:

Parameter name only The collected URLs display the parameter name only and remove the parameter's value.

None The parameter's name and value are displayed for the collected URLs.

5. Click **Add** to add a rule exception or **Edit** to edit an existing rule exception.
6. In the Add Rule dialog box, define the exception rule as follows:

When encountering parameter: Enter the name of the parameter you want to enforce the exception rule on.

Show: Select Parameter name and value, Parameter name only, or None to determine which information Precise for Web should show when it encounters the specified parameter.

7. Click **OK** to close the dialog box.

About configuring Display Settings

Selecting the option on the Display Settings dialog of the Settings menu allows you to maintain the selected tab or view when switching entities, to define the maximum number of items to display for each association, and the type of Web client activity. For example when the first option is selected and you are in the Activity tab Server-Side tab, selecting a different association will not change the Server-Side tab to Client-Side.

About configuring Tree View Settings

You can define the parameters controlling the Tree View display using the Tree View Settings dialog box. By using this dialog box you can define the Sort criteria of the Tree View and the maximum number of search results when using the Tree View search.

About configuring Time Frame Settings

You can determine the resolution of the data that is displayed in the overtime graphs using the Time Frame Settings dialog box. By using this dialog box you can define the default time frame to display.

Tasks common to most tabs

The following tasks are commonly performed in most tabs:

- [Switching to a different tab](#)
- [Selecting a time frame](#)
- [Selecting an entity](#)
- [Sending an email message](#)
- [Adding, viewing, and deleting Favorites](#)
- [Determining which table columns to display](#)
- [Copying data to the clipboard](#)
- [Exporting to the Precise Custom Portal](#)

Switching to a different tab

You can easily switch between the different tabs using the Tab Selection bar. When you start most Precise products, the Dashboard tab opens by default. For other Precise products, another tab will open by default. The button of the selected tab is displayed in orange.

To select a tab

- Click a button on the Tab Selection bar to display information on the selected entity in a different tab.

Selecting a time frame

You can configure Precise for Web to display transaction performance data for a specific time frame using the predefined time frame options or calendar icons.

Selecting a predefined time frame from the toolbar displays transaction performance data for the selected time period up to the current time.

See [Selecting a predefined time frame from the Precise for Web toolbar](#) on page 19.

Selecting the time frame using the calendar icon, you can choose to define a time range independent of the current time, or to define a time range up to the current time.

See [Selecting a time frame using the calendar icon](#) on page 19. The predefined time frame options are:

- Last 20 minutes (20m)
- Last 8 hours (8h) (default)
- Last 1 day (1d)
- Last 2 weeks (2w)
- Last 5 weeks (5w)

The time frame selected affects all information displayed in Precise for Web. Only data that falls within the selected time frame is shown in these areas.

Selecting a predefined time frame from the Precise for Web toolbar

To select a predefined time frame

- From the Precise for Web toolbar, select one of the predefined time frames.

Selecting a time frame using the calendar icon

To select a time frame

1. Click the calendar icon. In the dialog box that is displayed perform one of the following:
 - a. To define a time frame independent from the current time, select the 'Time Range' option and select the Start and End dates and times.
 - b. To define a time frame up to the current time, select the 'Last' option and enter the desired time frame.
 - c. To use one of the three previously used time frames, select the 'Recently used' option and from the drop-down menu select the desired time frame.
 - d. To use a previously saved time frame, select **Use a previously saved time frame** and from the drop-down menu select the desired time frame.
2. To save your settings for future access, select **Save these definitions for future use as:** and enter a name in the corresponding field.
3. Click **OK**.

Selecting an entity

You can select the entity that you want to examine in more detail. To select an entity

- On the Main area of the tab, select the entity you want to examine.

Sending an email message

You can send an email message to one or more recipients from the Precise toolbar. The default subject for the message will be "Link to a Precise application".

The email will include a link to the Precise product in the current context (time frame and selected entries).

To send an email message:

1. Click the email icon on the Precise toolbar. The default email program opens.
2. Fill in the required fields and click **Send**.

Adding, viewing, and deleting Favorites

The Favorites feature allows you to save a specific location in your application and to retrieve the same location later without having to navigate to it.

About the Favorites feature

The new Favorites feature includes the following options:

- Relative Time Frame – Saving a relative time frame instead of static date. For example, saving the last 7 days will always display the last 7 days, depending on the day entered.
- One click to specific location – Once you open Precise by launching a saved Favorite item, you will not have to enter a login credential nor click the login button.
- IE Favorites support – Adding a new Favorite item in Precise will also add it to the IE Favorites menu.
- Auto Complete – The Favorites dialog includes a new combo box which supports AutoComplete.
- Auto Naming – The Favorites dialog generates item names based on the current location.

UI description

An Add/Delete Favorites option under the Favorites menu allows you to save the current location or delete an existing one.

To add a new Favorite location

1. On the Add/Delete Favorites dialog box, enter the name of the new Favorites entry.
2. Click **Add**. The dialog box is closed and the new Favorite is added to the list.

To view a Favorites location

1. On the Precise bar, click **Favorites**.
2. Select the Favorites location you want to view.

To delete an existing Favorite location

1. On the Add/Delete Favorites dialog box, select the Favorite location to be deleted.
2. Click **Delete**. The dialog box closes and the selected Favorite is deleted from the list.

NOTE The favorite address is displayed in the Address field and cannot be edited.

Determining which table columns to display

Tables are used to display information about a set of related entities in the Main and Association areas. It is possible to determine which columns to display in the Association area tables.

To determine which columns to display in the Association area

1. Click the Table icon on the upper right-hand side of a table and select **Column Chooser**.
2. In the Table columns dialog box, click the arrows to move the names of the columns that you want to display to the Visible box and the ones that you do not want to display to the Invisible box.
3. Click **OK**.

Copying data to the clipboard

At times you may want to save data displayed in the table area in a Microsoft Excel spreadsheet for further analysis or save an image of a graph to the clipboard.

To copy data displayed in the Association area to the clipboard

- Click the Table icon on the upper right-hand side of a table and select **Copy to clipboard**.
The table can be pasted into Microsoft Excel or as an HTML file. To copy a graph to the clipboard
- Right-click a graph and choose **Copy to clipboard**.
You can now paste the image into any application that works with the clipboard.

Exporting to the Precise Custom Portal

The Export to the Precise Custom Portal Portlet feature enables you to export the view of the chosen table or graph and generate a portlet with that view in the Precise Custom Portal, so that it will provide you with another way of monitoring your application.

Prerequisites

To be able to use this feature, you need to have the following rights in Precise:

- View permissions to all Tiers in the application
If you do not have sufficient rights, you will get an error message when trying to execute this feature.

Exporting the information

You can either export a table view or a graph view.

NOTE The name field has the following restrictions: maximum 100 characters.

To export a table view

1. Click the Column Chooser icon.
2. Select Export to the Precise Custom Portal Portlet.
3. Insert a name in the name field that clearly describes the table view.
4. Click **OK**.

To export a graph view

1. Right-click the graph.
2. Select Export to the Precise Custom Portal Portlet.
3. Insert a name in the name field that clearly describes the graph view.
4. Click **OK**.

Getting an overview of your application

This section includes the following topics:

- [About the Dashboard tab](#)
- [How the Dashboard tab is structured](#)


About the Dashboard tab

The Precise for Web Dashboard tab provides a comprehensive overview of all instrumented instances, based on the data collected by the Precise for Web agent. The information displayed in this workspace pinpoints performance trends and issues in your Web server and offers clear navigational recommendations throughout for further analysis and handling. See [About the Activity workspace](#) on page 24. See [About the Statistics tab](#) on page 32.

How the Dashboard tab is structured

The Dashboard tab is divided into three areas: the Main area, the Association area, and the Associated Findings area. The Main area lists all the instances that belong to the Tier. The Association area and Associated Findings area provide more extensive information on the selected instance.

Figure 3-1 Dashboard tab



About the Main area in the Dashboard tab

The instance table displayed in the Main area lists all the Web instances that are monitored by Precise for Web agents in the selected Tier. Each row corresponds to an instance. The All row in the main instance table summarizes the availability, response time and SLA compliance for all listed Web instances.

The table below describes the information displayed in the Main area when client-side collection (with server-side collection) is configured.

Table 3-1 Client-side collection (with server-side collection)

| Column | Description |
|---------------------|---|
| Icon | Indicates whether or not the instance is available. A green running figure indicates that the instance is currently available; a red standing figure indicates that the instance is currently unavailable. A graph with a red X indicated that this is an unmonitored instance. |
| Instance | Name of the instance (as configured during the Installation procedure). |
| Server | Name of the server where the instance is installed. |
| Response Time (Avg) | Average response time for all pages. |
| Web SLA Compliance | The Web Service Level Agreement broken down into the following components: SLA Breach, SLA Near-Breach and SLA OK. |
| Executions | Total number of instance pages viewed within the selected time frame. |

About the Association area and Associated Findings area in the Dashboard tab

This area is divided into two sections. The top section shows an overtime graph of the average response time vs. executions. The bottom section shows the findings. Clicking on the hyperlinked finding title or following step of the finding launches, in context, to the related tab from which the data was derived.

Examining performance over time

This section includes the following topics:

- [About the Activity workspace](#)
- [How the Activity workspace is structured](#)

About the Activity workspace

The Precise for Web Activity workspace displays an invocation tree of all monitored instances and various analysis tabs, enabling comprehensive and effective drilling down in the monitored instances and their invoked URLs to locate specific performance issues and their underlying causes.

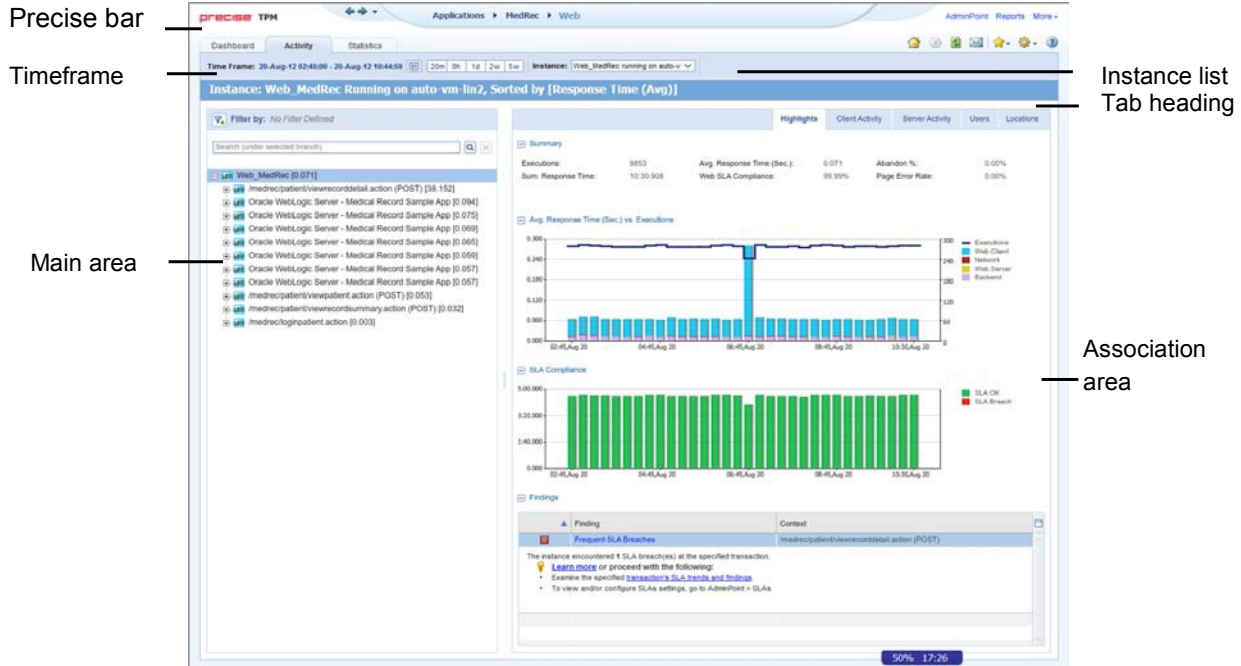
The Activity workspace allows you to analyze performance behavior over time. The information displayed in the Activity tab is derived from historical data and enables you to access information for different time intervals at varying levels of detail. You can analyze what happened at a specific time, pinpoint problematic time periods, and drill down to analyze a bottleneck that caused a performance problem.

The Activity tab was designed to answer questions such as:

- What happened yesterday, as well as a month ago or a year ago?
- What has been the performance level of my applications over time?
- Which activity suffered from poor performance? Which user was affected?
- What caused the poor performance? A problematic network? Poor page design?
- Should I analyze the performance of the internal back-office servers, such as the J2EE servers or the RDBMS, to find the root of the problem?

How the Activity workspace is structured

Figure 4-1 Precise for Web Activity workspace structure



The Activity workspace is divided up into two main areas. The Main area (left pane) features an invocation tree of all monitored instances. The Association area (right pane) features analysis tabs that provide specific information about the node selected in the invocation tree in the left pane. The workspace heading displays the instance name, the server name, and the field the by which the tree is sorted.

About the Main area in the Activity workspace

The Main area shows a tree structure with information on the selected entity. You can select a branch of the tree to examine the entity from different angles. For example, if you are analyzing an Instance entity, you can focus on its URL in the sub-branch.

About the Association area and Associated Findings area in the Activity workspace

The Association area and Associated Findings area provide corresponding information on the selected entity in the tree. The information that is shown depends on the chosen entity:

Table 4-1 Chosen entity overview

| Entity | Tabs | Information |
|----------|-----------------|--|
| Instance | Highlights | Summary, Average Response Time vs. Executions graph, SLA Compliance graph, Findings. |
| | Client Activity | Heaviest Web Client Transactions table, Average Response Time vs. Executions graph, Users table. |
| | Server Activity | Content Types by Response Time table, URLs by Response Time table, Average Response Time vs. Executions graph. |
| | Users | Users table, Average Response Time vs. Executions graph, Transactions Executed by Selected User table. |
| | Locations | Locations table. |

| | | |
|-----|------------|--|
| URL | Highlights | Summary, Average Response Time vs. Executions graph, SLA Compliance graph, Findings. |
| | Users | Users table, Average Response Time vs. Executions graph, Transactions Executed by Selected User table. |
| | Locations | Locations table. |

In the Association area, you can also drill down to another entity by clicking a table row. A drill down affects all tabs. You can also learn more about a finding (online help will open) or proceed with the steps provided in the finding (you will be taken to the relevant window).

How the Activity workspace is structured

The Activity workspace analysis tabs display in-depth information about and in context of the node selected in the invocation tree, shown within the tree in bold for referencing. While navigating through the analysis tabs, you can select different hyperlinked entities associated with the originally selected node. When you select a hyperlinked node in an analysis tab it selects that node within the invocation tree, and all information displayed in the analysis tabs are refreshed to correspond to the newly selected node.

The following table describes the analysis tabs available in the Association area of the Activity workspace and the node types for which each tab is displayed. As you navigate through the invocation tree, the available analysis tabs change according to the selected node.

Table 4-2 Analysis tabs available in the Precise for Web Activity workspace

| Tab Name | Description | Displayed For |
|-----------------|---|---|
| Highlights | Displays a comprehensive performance overview of the node selected in the invocation tree. This view reveals performance problems of the entire call tree invoked under the selected node. | <ul style="list-style-type: none"> All entities |
| Client Activity | <p>Displays Web Pages related data:</p> <ol style="list-style-type: none"> Heaviest Web Client Transactions according to: <ul style="list-style-type: none"> Response time Executions Page (Transaction) size Number of server URLs (components) Abandon rate - pages that were aborted by the user before they finished loading Average Response Time vs. Executions graph Users that executed the selected transaction <p>Displays the First Byte time - the time it takes from the moment the user asks for a new Web page until the first byte arrives back from the Web server.</p> <p>Displays the Rendering time - The time it takes for a Web page to be loaded from the moment the first byte arrives until the Web page is fully loaded, or the user interrupts or abandons the loading process.</p> | <ul style="list-style-type: none"> Instances Web Client |
| Server Activity | <p>Designed for the Web server administrator.</p> <p>Displays a breakdown of activity according to content type (Images, JavaScripts, style sheets, etc.).</p> <p>For each content type it shows volume of data,</p> | <ul style="list-style-type: none"> Instances Web Client |

| Tab Name | Description | Displayed For |
|-----------|---|--|
| | <p>errors and cache % and the underlying URLs.</p> <p>Displays The time spent on network activity from the server side perspective. This includes the time it takes to read the request from the network and the time required by the network to send the response back to the client.</p> <p>NOTE: If the Insight Savvy for Network data is not associated with the server side performance data, this counter will be 0.</p> | |
| Users | <p>Displays the Application users that were active in the corresponding time frame on the selected entity.</p> <p>Displays an overtime graph of Avg. Response time vs. Executions.</p> <p>Displays the Platforms and Browsers that the user used to run the selected entities.</p> <p>According to this information you can determine how your application runs on different browsers.</p> | <ul style="list-style-type: none"> • All entities |
| Locations | <p>Displays and compares performance data of all the Locations. (A Location is based on IPs and can be set from the adminpoint).</p> <p>Here you can see if there is a problem connecting to your application from a specific location.</p> | <ul style="list-style-type: none"> • All entities |

Examining Web instance statistics

This section includes the following topics:

- [About the Statistics tab](#)
- [How the Statistics tab is structured](#)
- [About the entities you can examine](#)
- [How the Statistics tab can help you identify performance problems](#)

About the Statistics tab

The Statistics tab displays statistical information on the behavior of your applications Web instances with respect to the Web server. This statistical information can be used to identify and analyze the cause of a performance problem, especially with respect to the server side. It can also help you make future tuning decisions regarding your Web servers.

NOTE This tab is only available if you have server-side collection installed.

How the Statistics tab is structured

The Statistics tab displays information on a specific Tier and its related instances, or on a specific instance and its related performance counters.

When you open the Statistics tab from another tab, you will either be focused on an individual instance, or on the entire Tier (if you selected the ALL row in the Dashboard). The tab is always launched at the instance (or Tier) level, even if the currently selected entity (in another tab) is not an instance or Tier. The selected entity is always reflected in the Tab heading, which serves as a point of orientation.

About the Main area in the Statistics tab

The Main area in the Statistics tab shows comprehensive information on the selected entity. You can choose from several views to examine the entity from different angles. You can for example focus on CPU Utilization, Load, or Memory.

About the Association area in the Statistics tab

The Association area in the Statistics tab provides information on the entities associated with the selected entity (displayed in the Main area). For example, when you view Tier information, the Association area displays the information for the instances belonging to that Tier.

From the Association area, you can also drill down to another entity by clicking a row in the table. A drill-down affects the whole tab.

When you drill down to another entity, the following are displayed:

- The Tab heading displays the new selection
- The Association area displays the statistics counters that were collected for the selected instance

The Association area displays all the Performance counters collected for the selected instance. The table displays the performance counter name and value. This list can differ to reflect the specific information relevant to your systems configuration type.

NOTE To gain a better understanding of the processes involved and better tune and improve your Web server's performance, see the counters explanation provided by your specific Web servers documentation.

How the Statistics tab can help you identify performance problems

From the Association area, you can click the row for the performance counter to view detailed information for the selected counter. The Tab heading indicates the newly selected entity, and the Main area displays an overtime graph for the performance counter you drilled down to.

About the entities you can examine

The Statistics tab displays information on different entities. This section provides an overview of all entities, their meaning, and views.

The following entities are available in the Statistics tab:

- [About the Tier entity](#)
- [About the Instance entity](#)
- [About the Performance Counter entity](#)

About the Tier entity

The Statistics tab, when launched from the Dashboard tab with the All row in the All Instances table selected, opens with information on the Tier level. The Tier entity displays information regarding the load of the Tier. For this entity, the Load view is the only available view.

The Load view displays the following graphs:

- Requests
It displays the amount of Web server requests logged within the selected time frame.
- Concurrent Connections
It displays the current number of connections to the service within the selected time frame. You can associate an Tier with instances.

The table below describes the information that is displayed in the table, in the Association area, for instances associated with the Tier.

Table 5-1 Association area table information

| Table Column | Description |
|------------------------|---|
| Instances | Name of an instance (as configured during the Installation procedure). |
| Server | Name of the server where the instance is installed. |
| Requests | Total number of HTTP requests logged for the instance within the selected time frame. |
| Concurrent Connections | Number of current connections to the service. This is the sum of both non-anonymous (authenticated) and anonymous (unauthenticated) users. If this number has reached or is approaching the maximum number of allowed connections, the Web service is running at full capacity. |
| Memory | Amount of virtual address space reserved by the Web server process. |
| Page Faults | Number of times a virtual page was not found in memory. |
| CPU Utilization | The time spent, in percentage, processing threads by all CPUs. |
| Cache Hit Ratio | The percentage of total requests that were serviced from the Web server output cache. |

NOTE These column headings are derived from counter names. To gain a better understanding of the processes involved and to better tune and improve your Web server's performance, see the counters explanation provided by your specific Web server's documentation.

About the Instance entity

You can view Instance level information by performing any of the following options:

- In the Association area of the Statistics tab, drill down to a specific instance by clicking the relevant table row.
- In another tab, drill down to a specific instance and click **Statistics** on the Precise Tab Selection bar.
- In the All Instances table of the Dashboard tab, select an instance and click **Statistics** on the Precise Tab Selection bar.

The Instance entity displays information on the most important performance counters. The Instance entity provides the following views:

- [About the overview for an instance](#)
- [About CPU utilization for an instance](#)
- [About instance memory information](#)
- [About instance load information](#)

The Association area displays all performance counters collected for the selected instance. The table displays the name and value for each performance counter. The listed counters can differ to reflect the specific information relevant to your systems configuration type.

NOTE To gain a better understanding of the processes involved and better tune and improve your Web server's performance, see the counters explanation provided by your specific Web servers documentation.

About the overview for an instance

In a bar graph, the Overview displays request counter data within the selected time frame and information regarding common statistical counters collected for the selected instance.

The table below describes the relevant counters

Table 5-2 Common statistical counters

| Counter | Description |
|-----------------|---|
| CPU Utilization | The percentage of time spent processing threads by all CPUs. |
| Memory | The amount of virtual address space reserved by the Web server process. |
| Page Faults | The number of times a virtual page was not found in memory. |
| Access Count | The total amount of HTTP requests logged for the instance within the selected time frame. |

Table 5-2 Common statistical counters

| Counter | Description |
|------------------------|---|
| Concurrent Connections | The number of current connections to the service. This is the sum of both non-anonymous (authenticated) and anonymous (unauthenticated) users. If this number has reached or is approaching the maximum number of connections allowed, the Web service is running at full capacity. |
| Cache Hit Ratio | The percentage of total requests that were serviced from the Web server output cache. |

About CPU utilization for an instance

In a bar graph, the CPU Utilization view displays CPU Utilization counter behavior for the selected instance within the selected time frame.

About instance memory information

The Memory view displays the following bar graphs:

Web Server Memory Memory counter behavior for the selected instance within the selected time frame.

Page Faults Page Faults counter behavior for the selected instance within the selected time frame.

About instance load information

The Load view displays the following bar graphs:

Access Count Requests counter behavior for the selected instance within the selected time frame.

Concurrent Connections Concurrent Connections counter behavior for the selected instance within the selected

time frame.

About the Performance Counter entity

In a bar graph, the Performance Counter entity displays the behavior of a particular performance counter over time, such as: TotalProcessTime, ProcessMemory, or PageFaultCount.

How the Statistics tab can help you identify performance problems

The Statistics tab provides many performance counters grouped into several predefined graphs that display related counters that enable you to locate performance problems in your system.

The View control drop-down menu, in the Main area, lets you choose different view options for your selections. You can identify a performance problem by doing one or more of the following:

- [Examining memory statistics](#)
- [Examining CPU utilization statistics](#)
- [Examining load](#)

Examining memory statistics

You can examine a Web instance within a selected time frame to confirm that memory usage falls within the normal range and that no processes are lacking or consuming too much memory. Lack of memory or a high memory consumption rate can cause paging. A page fault occurs when a process requires code or data that is not in its working set (its space in physical memory), thereby causing the page to be retrieved from the disk and written to the physical memory. If there is not enough space in the physical memory, a chosen page needs to be written to the disk to free space.

To examine memory statistics

1. In the Time Frame list, choose the period of time you want to analyze.
2. On the View controls in the Main area, click **Overview** and analyze the graph for requests to determine whether paging is the source of the performance problems.

If you conclude that excessive paging is the source of the performance problems, examine the amount of memory allocated to the Web server and reduce it, if necessary.

3. On the View controls in the Main area, click **Memory**.
4. In the Association area, drill down to the available pages and performance counters to analyze memory usage.

Examining CPU utilization statistics

You can examine a Web server instance within a selected time frame to confirm that CPU usage falls within the normal range. High CPU usage may indicate the need to enhance performance. Low CPU usage may indicate a shortage of processing power.

To examine CPU utilization statistics

1. In the Time Frame list, choose the period of time you want to analyze.
2. On the View controls in the Main area, click **Overview** and analyze the graph to determine whether CPU usage is the source of the performance problems.

If you conclude that high CPU usage is the source of the performance problems, consider adding more CPU or tuning the application in the Activity tab.

3. On the View controls in the Main area, click **CPU Utilization**.
Examine the graph to determine whether a high level of CPU utilization is causing processor bottlenecks.
4. On the Association controls in the Association area, click a counter and check the results.

Examining load

You can examine Web load activity and performance within a selected time frame. A high load may indicate the need to distribute the work activity to enhance performance.

To examine load

1. In the Time Frame list, choose the period of time you want to analyze.
2. On the View controls in the Main area, click **Overview** and analyze the graphs to determine whether excessive load is the source of the performance problems.
3. On the View controls in the Main area, click **Load**.
The Load view displays requests and concurrent connections. Compare these graphs to determine where most of the activity is being generated.
4. On the Association controls in the Association area, click a counter and check the results.

Web Findings

This section includes the following topics:

- [Frequent SLA Breaches](#)
- [Hotspot Detected](#)
- [Significant Backend Activity](#)

Frequent SLA Breaches

SLA thresholds are defined to help the user pinpoint transactions experiencing performance issues according to specific criteria. Frequent SLA breaches and near breaches can be caused by an underlying performance issue.

Working with the finding

To effectively locate the root cause of the performance finding, perform one of the following:

- Click on the transaction's link, and then look at the overtime SLA behavior to locate and zoom in to a specific (problematic) time frame. View the findings for that time frame and drill down until you locate the root cause.
- Select the root level of the transactions tree and select the Transactions tab. A high rate of SLA breaches across the application could result from overall resource exhaustion. Open the Tiers tab, follow the link to the heaviest ones and open the Statistics tab to learn more about the environment performance issues, like high memory usage, CPU usage and so on.
- Go to AdminPoint>Settings>SLA to view the thresholds definitions. (When you have too many SLA breaches, it may be a result of thresholds that are not defined appropriately for your application).

Hotspot Detected

A high work time for a specific hotspot (reflecting the hotspot's work time only, without the underlying call path), can indicate a performance issue within the context of that hotspot.

In the same way, a high work time for a specific occurrence of a hotspot invoked multiple times in the invocation tree can indicate a performance issue within the context of that specific occurrence.

Working with the finding

To effectively locate the root cause of the performance finding, perform the following:

- Examine the heaviest occurrences further by following the featured link. The invocation tree opens to the hotspot's heaviest call path, facilitating effective navigation to the root cause. Examine the information displayed and look at the overtime graph and findings to drill down further to find the root cause of the performance issue.

NOTE By default, information is displayed for the heaviest hotspot's call path. To investigate the other call paths, select them from the invocation tree. (They are highlighted in bold).

Significant Backend Activity

A significant percentage of the selected entity's total service time is spent waiting for back-end activity (for example, J2EE, .NET and so on).

A possible solution is to consider tuning your transaction performance and the time spent executing external activity in the appropriate tier.

Working with the finding

To effectively locate the root cause of the performance finding, perform the following:

- Select the link in the expanded finding area in order to view this transaction in the Transactions tab.
- Analyze the transaction in order to determine the appropriate back-end tier that should be investigated next.

| | |
|----------------------------------|---|
| abandonment rate | In Web, a counter that keeps track of the percentage of users that abandon the loading of the Web page before it completes downloading. |
| action | An operation that Alerts FocalPoint automatically performs when detecting a warning or critical status for a specific metric. According to the defined action, Alerts FocalPoint opens a message box, sends an email or SNMP trap, or executes a program. |
| action item (manual) | Some installation steps cannot be executed automatically by Precise Installer. The administrator is required to execute them manually. Action Items are presented (when necessary) at end of installation or product update. |
| AdminPoint | The central administration console of Precise that facilitates the maintenance, configuration, and management of all installed Precise components, such as monitoring the status of all Precise agents and PMDB processes, getting license information, starting and stopping the agents, getting log data on agents and events, changing PMDB settings, and installing patches. See also PMDB . See also agent . |
| advice | In Oracle, an algorithm that is designed to recommend on gathering statistics, creating new indexes, change existing indexes, and add or delete hints to make Oracle's Optimizer choose a better access path and make the statement perform better. Can be activated from any DML (Data Manipulation Language) statement. |
| agent | A program that runs on a server machine to collect, process, or communicate performance information. The Precise installation consists of multiple agents. |
| alert | The state of an Alerts metric that has reached a near-critical or critical status. An alert is issued by Alerts, triggering an action and informing of a problem that has occurred or is likely to occur within the area sampled by the specific metric. See also action . |
| Alerts | These product provide alerts to problematic conditions before they turn into performance problems, based on predefined metrics and thresholds. Alerts can automatically perform an action, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program. |
| Alerts FocalPoint | An agent that receives data from the InformPoint agents, stores it, and performs any action that has been user-defined for that specific alert, such as displaying a pop-up message, sending an email message or SNMP trap, or running a program. See also InformPoint . |
| alert type | In Alerts, the status of all metrics belonging to a metric group or a monitored instance, indicating the current performance level through colors. See also alert . See also metric . |
| application server metric | In J2EE, a metrics that is provided by the application server or by customer code. This can include metrics published by the Java Management Extension (JMX) APIs or vendor-specific APIs, such as IBM's Performance Management Interface (PMI). |
| AppTier | The abbreviation for an application tier in a Precise environment. An AppTier contains one or more instances of the same technology and purpose. Application tiers do not necessarily correspond to distinct physical servers: in many cases, the tiers are logical, with one server running multiple AppTiers or one AppTier spans a cluster of servers. A Precise environment may contain multiple AppTiers on the same technology. For example, you may group J2EE instances (JVMs) into a J2EE Presentation AppTier and a J2EE Business Logic AppTier. Segmenting application service time into the contribution of individual application tiers is helpful in identifying the source of performance problems. Analyzing the performance and behavior of each tier separately is crucial for isolating the root causes of performance problems. |
| cabinet | In Oracle, SQL Server, DB2, and Sybase, the highest logical level in the SQL workspace hierarchy. A cabinet contains folders and, within folders, statements. |

| | |
|-------------------------------|---|
| call path | A subset of a stack trace including only those methods that have been instrumented. |
| client cache | In Web, a counter that keeps track of the percentage of requests taken from the client cache (http status 304). |
| client-side collection | See Web client browser-side . |
| cluster | A group of servers or instances that are configured to be treated as a single entity. In Oracle, all Oracle instances of the same Oracle database (RAC configuration) are treated as a cluster and presented in the Precise for Oracle UI as a Database. In J2EE, a group of JVMs that serve the same application can be defined as Cluster (even if they are not defined as a cluster in the application server configuration). All JVMs of the same cluster share the same configuration. |
| collapsed access plan | In SQL Server and Sybase, the access plan of a unique group of statements or batches that belong to the same collapsed statement or batch but have different access plans. This can differ due to the constants in the text of the original (not collapsed) statements. See also explained statement . See also collapsed statement . |
| collapsed statement | In Oracle, SQL Server, Sybase, and DB2, a statement whose constants are replaced with parameters. Each collapsed statement can have several access plans, according to the occurrences of its statements. See also collapsed access plan . |
| Collector | The program that runs on the monitored server to collect performance information. Some technologies allow a single Collector (agent) to serve multiple instances running on the same server. Other technologies require a dedicated Collector per monitored instance. In most cases, it must be installed on the monitored server. Collectors for SAP, SQL Server, and Sybase may reside on a remote server. See also agent . |
| collector slice size | In J2EE, the smallest unit of data presentation and storage. It can be configured as 30 seconds or five minutes. If SmartLink is configured, relevant Web and Microsoft .NET Collectors will report to SmartLink using the same collector slice size. Regardless of the defined collector slice size, the most granular slice size presented on both the Precise for Web UI and the Precise for Microsoft .NET UI will be 15 minute intervals. |
| completion rate | In Web, a counter that keeps track of the percentage of users that completed the loading of a Web page. |
| count | The number of occurrences observed during a measurement interval. |
| CPU time | The average amount of time consumed by the operating system actively processing instructions on behalf of a running activity. |
| critical status | In Alerts, the status represented by a red bullet indicating that the value of the sampled metric has exceeded the near-critical and critical threshold values. See also metric . |
| Cross-AppTiers | In Insight, Report Manager, and Alerts, a perspective that provides high-level information about the performance of all AppTiers in an environment, including operating system data. It provides a holistic view of the entire environment and helps understand how the AppTiers interact. See also AppTier . |
| current data | In SmartLink, the end-to-end activity, detected by Precise over the last five minutes, including activities that are still in progress. In all databases products, a workspace that presents Current Sessions and the SQL statements that were executed including statements that are still in progress. |
| custom | Represents invocations of a type not specifically defined. See also work time . |
| customized metric | In Alerts, a user-defined metric that measures site-specific parameters. |
| drilldown | Within Precise products and Insight SmartLink, the filtering of analyzed data by clicking a specific entity. Then, additional information about the selected entity is presented, plus a breakdown of its activity by another entity. For example, drilldown on a transaction and get an overtime graph of the selected transaction, plus a list of top users that executed the selected transaction. |
| entry point | Usually a top-level HTTP or EJB request. An entry point can originate when a user clicks in a browser or an E-commerce server invokes a remote EJB. J2EE makes a distinction between service request invocations and other invocations. The first HTTP or EJB invocation within a call path is designated as an entry point. |

| | |
|----------------------------|--|
| environment | <p>The highest Precise logical group. It may contain multiple AppTiers of various technologies that serve an application together. For example, a Payroll Production environment may contain all Web servers, application servers, transaction managers, databases, and servers that serve this application. Alternatively, it may contain any set of instances that form together an administrative group. Since Precise version 8.5, an instance can belong to more than one environment. Let's say a single Oracle database serves two different applications plus the DBA wants to associate this database with a group of other databases under his responsibility regardless of served applications. In this case the Oracle database will be associated to three different Precise environments.</p> <p>See also AppTier.</p> <p>See also instance.</p> |
| ERP Extension | <p>In Oracle, DB2, and SQL Server, an extension to the Collector that provides detailed information on the activities and resource consumption of packaged application components. It correlates the database information, and the packaged application information and lets you see users, transactions, reports, and other elements of ERP applications, such as Oracle Applications, SAP, PeopleSoft, and Siebel.</p> |
| event | <p>In Alerts, the occurrence of a sampling or progress. A sampling occurrence occurs every time a metric samples. A progress occurrence occurs when a metric's progress status is changed or when the investigated status reaches the end of the given investigation time. In AdminPoint, all occurrences reported by Precise agents, including informational events, warnings, and errors related to one of the agents. All events are shown in the Events view.</p> <p>See also metric.</p> <p>See also progress.</p> |
| executions | <p>In Oracle, SQL Server, DB2, and Sybase, the number of times a SQL statement was executed during the selected time frame. In SmartLink and SAP, the number of times a transaction was executed during the selected time frame.</p> |
| explained statement | <p>In Oracle, SQL Server, DB2, and Sybase, a statement whose access path (chosen by the RDBMS Optimizer) is clarified and translated into a visual display. Explained results include information on the objects referenced by the statement and the operations performed on these objects.</p> |
| extended collection | <p>In Oracle, a function that proactively specifies a future period during which Oracle activity data is collected and organized for subsequent analysis. Extended collections are an easy means to view collected information, assess application resource consumption, and identify bottlenecks that are inhibiting application performance and end-user productivity.</p> |
| Federated Precise | <p>Federated Precise is a version of Precise (starting with version 8.5) that can manage multiple Precise installations within unified StartPoint and AdminPoint screens, displaying and managing all environments, instances, and installations.</p> |
| findings | <p>A ranked list of top performance problems in the selected context. They are presented in Oracle, SQL Server, .NET, and J2EE. Each finding provides: a short explanation of the problem, detailed background information, and most important: one or more links to further investigate the problem and possibly get more granular findings on the selected context.</p> |
| first byte time | <p>In Web, a counter that keeps track of the time that it takes from the moment a new Web page is called until the first byte arrives back from the Web server.</p> |
| FocalPoint | <p>An agent that communicates with the Listeners installed on the monitored servers, receives data from the Collectors, periodically processes and stores this data in the PMDB, and serves UI requests.</p> |
| folder | <p>In SQL Server, Oracle, DB2, and Sybase, the intermediate logical level in the SQL workspace hierarchy. Folders are grouped into cabinets and contain SQL statements.</p> <p>See also cabinet.</p> |
| Framework Installer | <p>The application that facilitates the installation of Precise framework components. It can be invoked from the installation DVD to install a new Precise deployment. It can also be invoked from an existing Precise deployment to install an additional framework node and attach it to the originating Precise deployment.</p> <p>See also Framework node.</p> |
| Framework node | <p>A set of FocalPoint agents that are installed together and manages a set of monitored instances. The performance data of these instances will be loaded into a dedicated PMDB. A single Precise deployment may contain multiple framework nodes (using a separate PMDB for each node). An environment cannot span over multiple framework nodes.</p> |

| | |
|---------------------------------------|--|
| garbage collection | An automatic process in the Java runtime environment that periodically reclaims memory used by objects that are no longer referenced. The process can impact an application's performance while memory is being reclaimed. Java programmers may initiate garbage collection explicitly. |
| group | In Web, the identifier that is used to group other identifiers, such as sites or URLs. |
| hint | In Oracle, an instruction directed at the Oracle Optimizer that includes considerations for an execution plan. The Oracle Optimizer will build an execution plan based on the hint, ignoring its own set of considerations. |
| hour group | A unit that reflects the type and level of activity within the system at different times. By defining the times of the day that are peak and off-peak, or day and night, the performance analysis can be focused on those particular times of the day. If, for example, most performance problems occur within nighttime and weekend batches, it can be useful to focus only on them. |
| InformPoint | In Alerts, an agent that retrieves performance data from all installed Precise products, analyzes it, and sends an alert to the Alerts FocalPoint when the predefined thresholds are exceeded. See also agent . |
| Insight | The Precise product family that facilitates the process of monitoring and correlating system performance. It consists of Insight. See also Report Manager . See also FocalPoint . |
| Insight FocalPoint | An agent that receives performance information from Insight Savvies, which monitor the environment. Insight FocalPoint then correlates, processes, and stores this information in a centralized location. The Insight performance history is stored in the PMDB. See also Savvy . See also PMDB . |
| instance | A monitored object of a specific technology. The following list specifies what constitutes an instance for the various supported technologies: J2EE - a Java Virtual Machine (JVM - a logical name set by the user), Microsoft .NET - a Common Language Runtime (CLR - a logical name set by the user), Oracle Applications - an Oracle Applications Form server, Oracle - an Oracle instance, SAP - a SAP system, SQL Server - an SQL instance, Sybase - a Sybase instance, Tuxedo - a Tuxedo domain, Web - a Web server, WebSphere MQ - an IBM WebSphere Queue Manager, DB2 - a DB2 non-partitioned database or a DB2 database partition. During installation, the instance is associated with one AppTier and environment. An Instance can be moved to a different environment or associated with multiple environments without re-installation nor losing historical data. See also environment . |
| instance statistics | See performance counter . |
| instance/database changes | In SQL Server and Oracle, the part of the Collect Schema Changes process that captures instance definition changes and database option changes and saves them in the PMDB. |
| instrumentation | The process of inserting fault-tolerant recording hooks in Java byte code, .NET MSIL, HTML pages, or other monitored components, resulting in the capture of performance metrics. In J2EE, a mechanism that enables collecting performance information when an application is executed. The process involves inserting special fault-tolerant recording hooks into application class objects. In WEB, the insertion of recording hooks into HTML pages can either be in memory (Dynamic Instrumentation) or file-based (Static Instrumentation). |
| instrumentation context | See invocation context . |
| internal invocation | The process of invoking a request from an HTTP request (Servlet or JSP) or an EJB. J2EE displays an internal invocation like any other invocation. |
| invocation | An execution of a J2EE entity (a Servlet, EJB, SQL Statement, method, and so on). When mentioned in plural (Invocations), means the amount of times that the entity was executed. |
| invocation context | The context within which a method is invoked. For example, if Method A is invoked both from Method B and from Method C, there will be two different invocation contexts for the performance metrics collected for Method A, one for when it is invoked from Method B and one for when it is invoked from Method C. See also call path and instrumentation context . |
| JRE (Java Runtime Environment) | As the runtime part of the Java software, the combination of the components that enable the execution of a Java program: a Java virtual machine, the core class libraries, and the files that form the Java platform. |

| | |
|--------------------------------------|--|
| JSP (Java Server Page) | An HTML page with special tags for Java scripting. An application server processes the tags and generates a Servlet. |
| JVM (Java Virtual Machine) | An instance of a JRE that executes Java programs. A server-side Java application server is itself a Java program that runs inside a JVM. Servlets, JSPs and EJBs are Java programs (applications) that run within the application server's JVM. J2EE monitors the JVM running the application server and the server-side Java applications within the application server. |
| JVM Heap Memory key metric | The amount of real computer memory that is allocated to the JVM for executing Java programs. In Alerts, a parameter that monitors a very important performance aspect. The status of the Oracle instance (up/down), for example, is crucial for system performance. If the instance is down, this is the first problem that needs to be solved. Marking a metric as key metric ensures that a critical alert raised for this metric receives top priority by the person that is responsible to handle alerts at any time. It can also help determine which alert to handle first in case of multiple alerts. In Alerts, key metrics are always displayed at the top of a metric table when it is sorted by alerts so that they get immediate attention. |
| Listener | The agent that facilitates the communication between the various Precise agents across different servers must be installed on every server where Collectors or FocalPoint agents are installed. The Listener allows communication with all other agents installed on the monitored server, while only the Precise Listener port is known by other servers. |
| login name | In SQL Server and Sybase, the session identifier that represents the credential used to connect to the database. When an ERP extension is installed, the user name of the packaged application's client overrides the login name. For example, when SAP extension is installed, the SAP user name overrides the login name. |
| machine | A session identifier. A machine as sampled, for example, by the SQL Server Collector is the identification of the machine where the client process executes. Machine is also sampled by Oracle, DB2 and Sybase. In Insight terminology Machine is called Client machine. |
| Main framework node | A main framework node is the single point for login and also serves as the Precise FocalPoint for the entire deployment. |
| major collection count | Number of estimated major garbage collection events that occurred during the last J2EE collector slice. A major garbage collection can stop the application while JVM heap memory is being reclaimed. See also garbage collection . |
| major garbage collection time | Percentage of time spent by the JVM executing major garbage collection events during the displayed interval. See also garbage collection . |
| MBean | A Java object that represents a manageable resource. In J2EE, MBeans, or Managed Beans, are used for application server metrics. |
| memory logger interval | The interval at which the J2EE Collector gathers JVM heap memory data. All snapshots of the memory logger's data collected according to the memory logger interval are summarized with counters in the current aggregation interval. The memory logger interval's time span is typically a small fraction of the aggregation interval. |
| metric | In Alerts, a query that helps measure performance in the environment. Three types of metrics are available: System metrics relate to the internal resources, operations, and objects of the monitored infrastructure; application metrics reflect the way the applications perform; user-defined metrics can be customized to specifically relate to a site. When a metric's value exceeds one of the defined thresholds, its status changes to near-critical (yellow bullet) or critical (red bullet). See also key metric . |
| metric set | In Alerts, a unit that groups metrics that measure related performance aspects. The following metric sets exist: Status: includes metrics that alert to functional problems related to the instance. Performance: includes metrics that alert to performance problems related to the instance. Load: includes metrics that alert to instance-related load problems that may later cause errors or crashes. Service: includes metrics that alert to instance-related SLA breaches. The metrics in this set are sampled by Insight. Performance Trending: includes metrics that alert to potential future performance problems. The metrics in this set are sampled by Report Manager. |

| | |
|---------------------------------|---|
| | <p>Load Trending: includes metrics that alert to potential future load problems. The metrics in this set are sampled by Report Manager.</p> <p>Customized: includes user-defined metrics.</p> <p>Precise Status: includes metrics that alert to the near-critical or critical status of the installed Precise environment.</p> |
| module | A session identifier. In Oracle it contains the value of MODULE column in V\$SESSION table. In DB2 it contains the command name of non-SQL statements or the package name for SQL statements. |
| near-critical status | In Alerts, a status indicating that the value of the sampled metric has exceeded the defined near-critical threshold. A near-critical status is indicated by a yellow bullet. |
| network time | In Web, a counter that keeps track of the time spent on network activity from the server side perspective. This includes the time to read the request from the network and the network time to send the response back to the client. |
| node | <p>See Framework node.</p> <p>Note: when creating a new installation from DVD, it would be called "Framework." In the UI screens (columns in tables, choosing a system for a new environment etc), we would call them "Nodes." For example: choose a node for the new environment.</p> <p>When adding a new system within AdminPoint, it would be called a "Framework node".</p> |
| packaged application | <p>An application that is created and/or maintained by a third party and is not custom-built to one's specific needs. The following packaged applications have special treatment by Precise: SAP, Oracle Applications, PeopleSoft, and Siebel. They are harder to modify because the application code is either not available or hard to understand.</p> <p>See also ERP Extension.</p> |
| page size | In Web, a counter that keeps track of the amount of data loaded from the server to display the page. |
| page views | In Web, a counter that keeps track of the number of Web pages viewed at a specific Web site during a selected time period. |
| parent metric | <p>In Alerts, a joining of several child metrics. Each time a parent metric samples, it gathers data from a set of child metrics and presents it as a single metric query. The individual child metric values are displayed on the Thresholds tab of the Properties dialog box in Alerts.</p> <p>See also submetric.</p> |
| performance counter | In SQL Server, a Windows performance counter as reported by the operating system. In Oracle, an instance level statistic as gathered from some Oracle V\$ views. In Sybase, a sysmonitors counter collected by Sybase. In Web, an operating system or a Web server performance counter. |
| PMDB | The Precise data warehouse of performance and availability data. It can be hosted on an Oracle or SQL Server database. |
| portal server | An application server for Web-based applications that commonly provide personalization, single sign-on, and content aggregation from different sources and that host the presentation layer of Information Systems. Aggregation is the action of integrating content from different sources within a Web page. A portal may have sophisticated personalization features to provide customized content to users. Portal pages may have a different set of portlets creating content for different users. |
| portlet | A Java-based Web component, managed by a portlet container, that processes requests and generates dynamic content. Portlets are used by portals as pluggable user interface components that provide a presentation layer to Information Systems. |
| Precise deployment | An independent Precise system. It contains and manages various agent types and provides centralized monitoring and administration. A Precise deployment may contain multiple framework nodes. One of them is defined as the main framework node and it manages all other nodes of the Precise deployment. |
| Precise FocalPoint | See Precise FocalPoint agent . |
| Precise FocalPoint agent | An agent that manages all agents in a single Precise deployment. Additional product FocalPoint agents manage specific technologies and resources. |
| Precise Proxy agent | When multiple Framework Nodes are managed by a single Precise deployment, the Precise FocalPoint agent manages all of the agents of the main Framework node, while every other framework node has a Precise Proxy agent to manage all its agents. |
| Precise_root | This is the term used in a path for the Precise installation directory. The terms <code><i3_root></code> or <code>i3 root</code> can appear in text too. |

| | |
|----------------------------------|---|
| program | <p>A session identifier in Oracle, DB2, SQL Server, Sybase, Insight, and Alerts. A program as sampled, for example, by the SQL Server Collector, is the name of the executable that connects to the database. Applications that do not set the application name have N/A as program. When a SAP extension is installed, the SAP transaction overrides the program, and it may change during the application's lifetime. When a Siebel extension is installed, the Siebel views override the program, and it may change during the application's lifetime. In Insight, program is the name of an executable as recognized by the operating system. If an executable is invoked from a script (a batch or a shell), the script is displayed as the command entity. In Alerts, program is part of the customized metrics definition and is the name of the executable or stored procedure executed in the database that runs when the metric is sampled. In action definitions, program is the name of the executable that will run if the metric exceeds its threshold.</p> |
| progress | <p>In Alerts, the management state of a metric for which an alert has been raised. The following statuses exist:</p> <p>Open: An alert is raised.</p> <p>Investigated: The alert is taken care of.</p> <p>Closed: The problem has been solved.</p> |
| Proxy FocalPoint | See Precise Proxy agent . |
| recommended index | <p>In SQL Server, a function that uses the Microsoft Index Tuning Wizard to recommend on adding indexes or statistics for a selected statement, batch, or table. Based on the results of this function, the Optimizer will choose a better access plan and make the respective statement or batch perform better. For statements and batches, recommendations are based on the content of the statement or batch. For tables, recommendations are based on all the statements that are stored in the PMDB, are executed during the selected time period, and have an average duration time greater than the value defined in the registry (where 0 is the default).</p> <p>See also advice.</p> |
| related SQL | <p>In Oracle, a generated statement that uses alternative syntax to access the database in different ways and returns the same output as the original statement.</p> |
| relative frequency | <p>Number of contributor invocations per entry point. For example, if a service request calls three methods each time it is invoked, the Relative Frequency for the method is three invocations per service request execution. Similarly, if a service request calls one method every other time it is invoked (half of the time), the Relative Frequency for the method is 0.5 invocations per service request execution.</p> |
| rendering time | <p>In Web, a counter that keeps track of the time that it takes for a Web page to be loaded from the moment the first byte arrives until the Web page is fully loaded or the user interrupts or abandons the loading process.</p> |
| report | <p>A collection of queries, programming code, and layout settings that Report Manager executes to generate graphical results like tables and charts.</p> |
| Report Manager | <p>Uses historical information to identify problematic conditions, track long-term performance, volume trends and patterns, view availability problems over long periods of time and on different levels, compare the performance of similar systems, correlate between performance metrics of different products, assist in capacity planning, and generate demand-driven, user-defined reports.</p> |
| Report Manager FocalPoint | <p>An agent that examines the PMDB tables to produce scheduled performance reports.</p> <p>See also PMDB.</p> |
| report parameter | <p>In Report Manager, a keyword used in a report. Its value is set during the report execution. A parameter's value can be updated either permanently or for the current execution only.</p> |
| report property | <p>In Report Manager, the attributes that define a specific report, consisting of report parameters and scheduling information.</p> <p>See also report parameter.</p> |
| request | <p>In Web, a counter that keeps track of the number of HTTP requests sent for a viewed entity. request</p> |
| error | <p>In Web, a counter that keeps track of the percentage of requests completed with an HTTP error.</p> |
| sampling | <p>In Alerts, the process during which a metric queries a Precise product for a specific instance, retrieves values, and calculates the metric's alert level.</p> |
| sampling base | <p>In Alerts, the start time of a scheduled sampling process (by default Sunday, 00:00 AM).</p> |
| sampling period | <p>In specific Alerts metrics, the time frame during which statistical data is returned from other Precise products. Such a metric is for example Oracle's General Behavior metric, which returns database behavior for a certain period of time.</p> |

| | |
|--------------------------------------|--|
| sampling rate | In Alerts, the frequency of a metric's regular sampling schedule. The sampling rate is measured in minutes. |
| Savvy | An Insight agent that collects AppTier-specific performance data. |
| scalability | A system's ability to withstand load. For example, positive scalability means that the system continues to function properly even when it is called upon to service a larger number of users. |
| schema changes | In Oracle and SQL Server, a process that captures schema changes and saves them in the PMDB. In SQL Server, instance configuration parameters and database option changes are also captured. |
| Script Installer | A component of Precise Agent Installer that enables adding a small JavaScript script, also known as a Web browser-side agent, to the Web pages of your Web site. |
| server | Also called a host machine or server machine. The combination of a computer and associated operating system software that is accessed by multiple clients to perform requested functions and typically returns results to the clients. |
| server-side collection | In Web, the collecting of performance data from the Web server instance through the Web server agent. |
| service + network time | In Web, a counter that keeps track of the service and network time, including the total amount of time the request took to reach the server. |
| Service Level Agreement | See SLA (Service Level Agreement) . |
| service time | In Web and SmartLink, the time elapsed from when a request is received by the server to the time a response is sent back to the computer that made that request. The service time is measured on the server side. In J2EE, the time it takes an invocation to complete execution. In other words, service time is the average length of time between the start time and end time of a Java method execution. For example, the SQL service time is the time it takes the JDBC method executing the SQL statement to be completed. The service time includes CPU and wait time. The service times are reported in the interval in which they complete execution. Though a contributor's average service time may be very small, the contributor may be called many times. As a result, a contributor's overall contribution to performance may be large even though its average service time is low. |
| servlet | A class that is loaded only once and for which the application server uses multithreading to process requests. The servlet generates an HTML page that is sent back to the Web browser. |
| size | In Web, a counter that keeps track of the bytes sent and received. |
| SLA (Service Level Agreement) | A formal definition of an information system's performance goals. Within Precise, an SLA consists of clauses corresponding to various system activities. Once a system's SLA is defined, its SLA compliance can be analyzed, and breaches can be isolated to identify their causes. |
| SmartLink | In Insight, a component that provides transactive, correlated information across AppTiers from the user's perspective. Insight SmartLink functions in PeopleSoft and Web applications environments. |
| SmarTune | In Oracle, SQL Server, and .NET, an algorithm that analyzes performance metrics, identifies and ranks potential problem areas, and provides advice for correcting the problems. |
| socket | An I/O abstraction layer that enables processes to communicate with each other, regardless of whether they are on the same server. Sockets are bi-directional FIFO pipes that look like regular file I/O to the developer with the abstraction layer handling all of the low-level communication details. |
| StartPoint | The opening page of Precise. It provides a quick overview of the environment status and links to launch any of the Precise products. |
| Statistics workspace | In Oracle, DB2, SQL Server, Sybase, J2EE, and Web this workspace presents instance level statistics such as hit ratios, utilization, JMX metrics, logging, and more. |
| submetric | Also called child metric. In Alerts, the subquery of a parent metric. A child metric gathers its own data and combines it with the data gathered by other child metrics to form the result of the parent metric. Each child metric has its own thresholds and may be enabled or disabled individually. See also parent metric . |
| summary table | A container that stores the information collected by the Precise agents and loads it into the PMDB. Summary tables store the same data at different levels of granularity: time slice, hourly, daily, weekly, and monthly. By storing data in multiple summary tables, Precise can present a detailed view and progressively higher-level views of the same data. Summary tables are particularly useful for data aging. A data purging policy can be implemented for each summary table so that detailed data is retained for short-term historical analyses while more summarized data is used for long-term analyses and trending. |

| | |
|--|---|
| technology | A technology identifies the monitored object. For example, Oracle, SQL Server, Sybase, and DB2 are different technologies, while all Web servers (such as: Apache, IIS, and WebSphere) are defined as a single Web technology. A single monitored object can be monitored by two different technology's Collectors. For example, WebLogic server can be monitored by both a Web Collector and a J2EE Collector. |
| throughput | The number of average completions per second that are observed during an interval. |
| time slice | <p>A unit used to break up long sessions into smaller time periods. The length of a time slice is preset and cannot be changed. It represents the maximum time that passes before the data collected can be displayed. For example, if the length of a time slice is 15 minutes, the collection is only updated at 15-minute intervals.</p> <p>The length of a time slice is different for each technology: In Oracle, SQL Server, Sybase, DB2, Web, .NET, MQ, Tuxedo, Network, and OS, a time slice is 15 minutes. In J2EE and SmartLink, a time slice is either 30 seconds or 5 minutes. In Precise for SAP, a time slice is 5 minutes.</p> <p>See also collector slice size.</p> |
| URI (Uniform Resource Identifier) | The relative path to a resource after the location (network node) is found. |
| URL mapping | In Web, a function that defines rules that map URLs (Uniform Resource Locators) with dynamic parameters originating from a specific domain to a format that identifies the Web pages and prevents them from having different URLs. |
| Web client browser-side | In Web, the gathering of performance data from the browser running on the desktop of the Web application's user, through the static or dynamic instrumentation of Web pages. |
| Web server | A program that receives client requests for Web pages, retrieves the static pages and/or issues a request for dynamic page creation to an application server, and sends the pages to the client. In Web, a computer that delivers (serves up) Web pages. Every Web server has an IP address and possibly a domain name. |
| work time | The time spent in the specific entity excluding time spent by called (other) entities. In J2EE, work time of a method is calculated as the service time of that method minus the service time of the methods it invoked. |
| work type | In Web, a counter that keeps track of the type of the Web instance, such as Web, PeopleSoft, SAP, or Siebel. |
| workspace | A display unit in Precise products. All workspaces display data from different perspectives. For example, in Oracle, SQL Server, and Sybase, the Current workspace shows information on the sessions currently active in an application, and the Objects workspace displays information on Oracle or SQL Server database objects that can be used to understand relationships and associations between database schema objects. |